REVISION 1

REPORT DOCUMENTATION PAGE			*Form Approved* *OMB No. 074-0188*

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	18 March 2004	Symposium Paper 17-18 March 2004

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
Distributed Machine Intelligence for Automated Survivability	N/A
6. AUTHOR(S)	
Katherine F. Drew and David Scheidt	

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Office of Naval Research Hull, Mechanical and Electrical Division, Code 334 800 N. Quincy St. Arlington, VA 22217 Johns Hopkins University Applied Physics Laboratory NSTD/STI 11100 Johns Hopkins Road Laurel, MD 20723-6099	N/A

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING / MONITORING AGENCY REPORT NUMBER
	N/A

11. SUPPLEMENTARY NOTES

Prepared for the Engineering the Total Ship (ETS) 2004 Symposium held in Gaithersburg, Md. at the National Institute of Standards & Technology and sponsored by the Naval Sea Systems Command & the American Society of Naval Engineers

12a. DISTRIBUTION / AVAILABILITY STATEMENT	12b. DISTRIBUTION CODE
Approved for public release; Distribution unlimited	A

13. ABSTRACT *(Maximum 200 Words)* Future Naval platforms face new dynamic operational scenarios that demand more flexible performance. At the same time, reduced manning and lower total ownership costs are now major design and acquisition objectives. Improved warfighting capability can be achieved by reducing vulnerability to damage and failure events. Rapid system recovery from unanticipated damage using current doctrine and practice conflicts with today's reduced manning objectives. Decentralized ship system architectures and agent based technologies promise to enable the Navy to improve rapid system recovery and assist in meeting these affordability challenges. Decentralization of systems and resources improves both ship survivability and fight through capability. This is accomplished through rapid sensing and response and dynamic reconfiguration, which results in improved continuity of service of ship systems. Embedded intelligence at the component level insures rapid, effective autonomous reaction and response to local fault conditions. Agent based technologies are utilized to provide autonomous cooperation between sensors and actuators, in which elements reason and react locally while achieving global objectives through agent-to-agent communications. While intelligent decision making is performed locally by autonomous agents, the sailor will direct these agents through comprehensive supervisory control, with improved on demand situational awareness. When fielded, these systems will provide increased situational awareness, increased fight through capability, and improved damage control. This paper describes Navy Science and Technology projects currently underway in academia, industry and Navy laboratories to achieve these goals.

14. SUBJECT TERMS	15. NUMBER OF PAGES
	9
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	UNCLASSIFIED	UNCLASSIFIED	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

20040419 132

Distributed Machine Intelligence for Automated Survivability

Katherine Drew, Office of Naval Research
David Scheidt, Johns Hopkins University Applied Physics Laboratory

Abstract **- Future Naval platforms face new dynamic operational scenarios that demand more flexible performance. At the same time, reduced manning and lower total ownership costs are now major design and acquisition objectives. Improved warfighting capability can be achieved by reducing vulnerability to damage and failure events. Rapid system recovery from unanticipated *damage using current doctrine and practice conflicts with* today's reduced manning objectives. Decentralized ship system architectures and agent based technologies promise to enable the Navy to improve rapid system recovery and assist in meeting these affordability challenges. Decentralization of systems and resources improves both ship survivability and fight through capability. This is accomplished through rapid sensing and response as well as dynamic reconfiguration, which results in improved continuity of service of ship systems. Embedded intelligence at the component level insures rapid, effective autonomous reaction and response to local fault conditions. Agent based technologies are utilized to provide autonomous cooperation between sensors and actuators, where elements reason and react locally while achieving global objectives through agent-to-agent communications. While intelligent decision-making is performed locally by autonomous agents, the sailor will direct these agents through comprehensive supervisory control, with improved on-demand situational awareness. When fielded, these systems will provide increased situational awareness, increased fight through capability, and improved damage control. This paper describes Navy Science and Technology projects currently underway in academia, industry and Navy laboratories to achieve these goals.**

I. INTRODUCTION

Present day Naval warfare scenarios point to the need for rapid and flexible response to a broad spectrum of threats. Incidents such as the attacks on the USS Cole, Stark and Princeton have demonstrated that threats to Naval platforms can arise quickly, unexpectedly, and have huge adverse consequences for personnel and platforms. Survivability of Naval platforms when operating under unpredictable and lethal circumstances is obviously a critical concern for the Navy. Survivability consists of the elements of susceptibility, vulnerability and recoverability. [1] Recoverability is defined as crew actions to reconfigure and restore systems to enable the ship to carry out its missions under damaged conditions. Recoverability has the elements of minimizing the time required for restoration of mobility, seaworthiness, and crucial ship systems, and the sustainment of warfighting capability (sometimes referred to as "fight through"). Survivability objectives, such as rapid location and containment of damage, rapid restoration of operational effectiveness, and maintenance of C^3 Operational Condition without Interruption, all point to the need for most timely intervention in the case of a damage event or critical system malfunction. [1] Given Naval requirements for reduced onboard manning, and the DoD drive to lower Total Ownership Costs while maintaining ship performance, development of *technologies to maintain ship readiness* and recoverability under conditions of reduced manning is critical.

II. BACKGROUND

Early ships were controlled entirely by human operators. Control mechanisms were manual switches and valves. System knowledge was obtained via human inspection of mechanical sensors. System reconfiguration required humans to turn switches and valves manually. System operation required continuous manipulation of system actuators by human operators. Coordination between components within a system was performed by human operators. Larger systems required more than one operator. Coordination between these large systems was achieved through operator dialogues. [Figure 1].

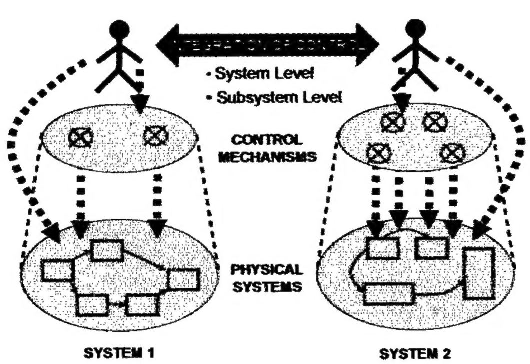

Figure 1 - Yesterday's Shipboard Control Architecture

The advent of automated closed-loop controllers reduced the frequency with which human operators

were required to manipulate individual components and subsystems consisting of small numbers of interconnected components. By enabling operators to consider a multi-component assembly as a single system, modern controllers enabled today's Navy to increase the complexity of ship control systems without corresponding increases in operator workload. The size of a system that can be controlled by today's controllers is limited. This limitation necessitates continuous direct human supervision of subsystems consisting of large numbers of components. Coordination between these large ship subsystems is still done sailor-to-sailor [Figure 2].

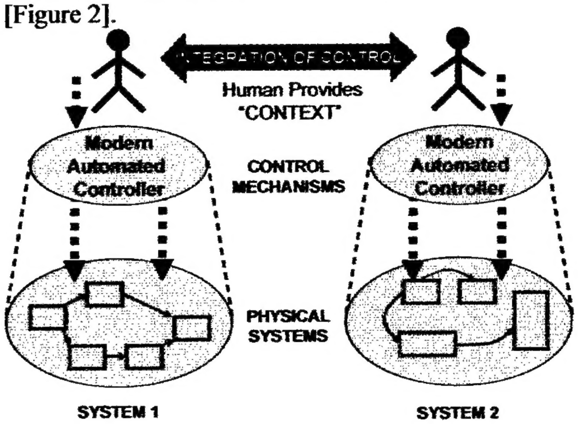

Figure 2 - Today's Shipboard Control Architecture

The complexity of ship systems is expected to increase exponentially with the advent of next generation ships. Continued reliance upon present-day automated controllers will necessitate a proportional increase in the workload required to operate ship systems. Fortunately, due to the growth of processing power available in microprocessors and advances in control theory, the sophistication of automated control continues to evolve. Individual devices now exhibit Component Level Intelligence (CLI), an example of which is the Smart Valve, which contains sensors, computation and communications capabilities. The Smart Valve can compute physical parameters such as flow rate and pressure associated with its operation as part of a shipboard fluid system,, and can share this information with its peers using a distributed control network.

The next step in the evolution of ship control will be the distribution of machine intelligence across multiple cooperating smart components. Initially, intelligent components will cooperate at the subsystem level. Distributed intelligence will enable subsystems to perform diagnosis, control and reconfiguration autonomously [Figure 3]. Example subsystems include valves, compressors, and power distribution modules. Above the autonomous system layer will be a system coordination layer, which will provide subsystem reconfiguration goals based on relationships between the machines in the subsystem level. The system coordination layer is concerned with optimizing the performance of the components and ensuring the availability of services. Examples of system functions at this system coordination level are Propulsion, Power Generation and Distribution, and Damage Control.

At the top, ship level control is concerned with ship wide resource allocation. At this level there are operator interfaces where the human can provide direction to the ship wide control system.

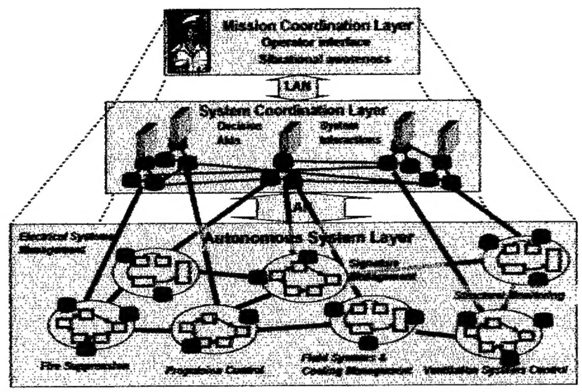

Figure 3 - Future Control Architecture

Allocation of responsibility between the operator and the automation systems will be a critical driver in determining future control system architectures. Ultimately, humans must be responsible for the operation of the ship. Even as the more tedious and dangerous tasks are off-loaded onto increasingly sophisticated autonomous controllers; the sailor must be made aware at all times of what the automated systems are doing. The control system must provide the sailor with situational awareness which, combined with operational context provided by the sailor, will enable shipboard systems to adapt in order to fulfill mission goals. Continuous feedback establishes a symbiotic

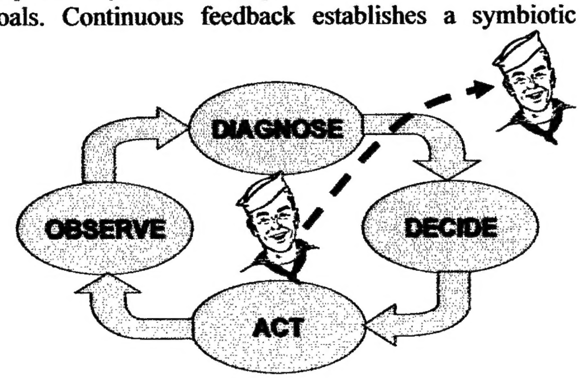

Figure 4 – Effects Based Control, Removing the Human From the Loop

relationship between the control system and the human operator. In this architecture, the role of the human will be to (1) establish mission objectives, (2) establish priorities for resource allocation, and (3) provide operational context. Humans will dictate goals to the system, and the system will determine what actions are to be taken to meet these goals and execute these actions. This method of operation is referred to as "effects based control", and is the goal of the architecture outlined here. [Figure 4] The use of agent based technologies and reasoning is a key enabler of effects based control.

III. AUTONOMOUS CONTROL

Autonomous (or automatic) control is closed-loop control that continuously reconfigures a physical system in response to feedback without requiring human intervention [2]. The reconfiguration triggered by the control system is guided by human inspired goals or objectives. We say human inspired rather than human-provided, because goals may be obtained either directly from humans during design, system start-up or run-time, or inferred from pre-existing goals by the control system. Regardless, the distinguishing feature of autonomous control is that it provides control continuously without halting operations to wait for human instructions.

Devising a policy for correct, complete control of a simple system may be accomplished by analyzing all possible states within the system through a failure means and effects analysis (FMEA). FMEA-based control policies consist of a set of predefined actions in response to system states. However, complete control policies derived from FMEA are only feasible for systems with few moving parts. As the size and complexity of a system increases, the number of possible system states increases exponentially, prohibiting the generation of a comprehensive policy in polynomial time. Systems that are small by Navy standards are far too complex to exert complete control through exhaustive rule sets. For example, a household system of twenty electrical circuit breakers with three states {on, off and tripped} has 3,486,784,401 possible states. Next generation auxiliary ship systems, such as those proposed for DD(X), are orders of magnitude more complex than household electronics and will consist of tens of thousands of connected components. The possible states of the DD(X)'s auxiliary systems will be on the order of 10^{5000}.

Control engineers manage complexity by predefining control policies for a small number of predefined nominal states and all possible single component failure within each nominal state. In our household electrical example two nominal states might be on and off. By considering single component failures only, the states that must be considered by our control policy are reduced from over three billion to a manageable twenty-two. This strategy for constructing a control policy works well when highly reliable components are used and component failure is limited to normal wear and tear. For example, using components with a mean-time to failure of less than 10^6 seconds will generate an unanticipated system state due to wear and tear roughly once every thirty-thousand years. Engineers account for these extremely rare unanticipated states by devising a control policy that moves the system into a universally safe state, usually shutting the system down.

These standard control practices cannot be relied upon to control Navy systems. First, the assumption that wear and tear is the primary cause for failure is false. Battle damage can be expected to cause multiple simultaneous failures. Second, fallback to a safe state when confronted with simultaneous failures is unacceptable. Continued functional performance in the presence of damaged components is a necessary requirement for all ship systems. Since comprehensive control policies for auxiliary ship systems are infeasible, and ship operations require that effective control be provided in the event of multiple simultaneous failures, correct control can only be achieved by establishing the control policy at run time. Automated mechanisms for devising control policies at run time are called intelligent control.

IV. INTELLIGENT CONTROL

Over the last fifty years a number of intelligent control techniques have been devised. All of these techniques involve *reasoning*, the ability to interpret information, usually provided by sensors, and generate knowledge through inference or deduction. Most reasoning techniques involve searching through possible future states and selecting those actions that generate the most advantageous futures as impetus for control. Others techniques, such as neural networks[3] and Q-learning[4], compare the effects of historical choices in order to continually improve a growing, emerging control policy. By themselves, these *learning* approaches are not suitable for auxiliary ship systems control as they require a tolerance for an occasional "wrong" answer from which the system learns[1].

Search-based approaches vary in both the strategy by which they search and by the method by which they

[1] Hybrid systems that incorporate learning approaches to reasoning as a component of the control system are a promising avenue of research.

generate future states for consideration. Early reasoning systems research represented future states as sets of unstructured facts [5]. Facts could be inferred by applying existing facts against a set of loosely organized rules. The successors to these early systems are Bayesian Belief Networks [6] which provide an ability to manage uncertainty and add structure to the application of rules. These rule-based systems have been shown to be effective at modeling the rules of thumb that constitutes human "expertise". However, the engineering effort required to extract the expertise represented by the rules used to generate a hypothesis is an arduous and inexact science. Further, the resultant rule set is difficult to validate other than through empirical testing which, due to the complexity of future Navy systems, is not feasible. Accordingly, these approaches to generating intelligence and their recent variants of Case-based Reasoning [7] and Partial Order Planning [8] are useful as operator assistants but may be inappropriate for autonomous control.

An attractive method for generating future states is the use of predictive system models that isomorphically represent the system being controlled. These methods of intelligent control typically rely upon qualitative models [9]. Qualitative models model physical components as separate logical units. How the model interacts with connected components is defined propositionally as a function of the component state. For example, the propagation of current across a simple electrical switch might be described by the following table [Figure 5].

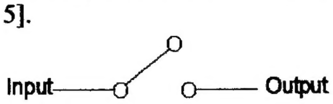

State	Output
Open	$i_{OUT} = 0$
Closed	$i_{OUT} = i_{IN}$
Shorted	$i_{OUT} <= i_{IN}$
Failed Open	$i_{OUT} = 0$

Figure 5 – Qualitative Model of a Switch

Models of large complex systems may be devised by modeling the relationships between component models. Component models, because they represent simple, well- understood systems, may be devised analytically and validated through empirical testing. Model validation may be assured by confirming that the structure of the model (component connections) accurately represents the physical system, thus resulting in a highly reliable method for system modeling. Sophistication, in terms of describing the continuous behavior of a component, may be described through hybrid models [10].

The search strategy used to select future states for investigation is a factor that determines the quality of the control strategy devised. Literally hundreds of search strategies have been developed over the last fifty years. However, the taxonomy of search algorithms may be organized by their approach to addressing three search related issues: goal or conflict resolution; breadth or depth first searching and stochastic or non-stochastic. A good general description of leading search algorithms may be found in [11].

V. SOFTWARE AGENTS

Intelligent control provides autonomous control the ability to address large complex systems. However, reliance upon a single intelligent controller presents a survivability problem. Damage to the centralized controller or to the communications infrastructure used by the central controller to communicate with actuators and sensors can result in a loss in controllability. Historically, the survivability risk of ship systems due to the loss of control was been minimal because intrinsically mobile human operators performed the task of high level control. Distributing control can be used to improve survivability as well as reducing the complexity of the system being controlled by an individual controller. Ideally, device controllers will be co-located with the device they are controlling decreasing the likelihood that an intact, serviceable device will be unable to function due to loss of control. In order to satisfy ship-wide goals distributed controllers must cooperate. Mintzberg [12] identified three cooperation mechanisms by which distributed controllers may coordinate: direct supervision, mutual

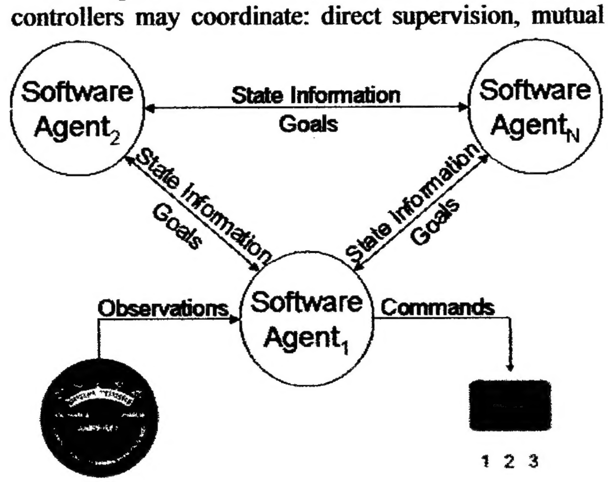

Figure 6 –Software Agents

adjustment and standardization. Standardization requires an a priori policy for decision making which,

as explained earlier, is not feasible for ship control systems. Distributed ship control systems will use either direct supervision, mutual adjustment, or both. The mechanism used to trigger distributed control will be *software agents*. Software agents are software processes that reason about and act upon their environment – in this case the devices being controlled [Figure 6]. The software agents used for future autonomous control will be intrinsically permanent and stationary, deliberatively diagnosing the state of and creating plans for physical components within their scope, in response to defined goals. Extrinsically, these software agents will be socially independent, requiring no actions on behalf of other agents yet sharing state information and goals with other agents in order to provide ship-wide control.

The cooperative mechanism used by agents to coordinate dictates the organization of the agent system. Supervisory control forms a basis for centralized or hierarchical control while mutual adjustment forms a basis for flat or heterarchical control. The simplest organization structure is centralized control. In centralized control, [Figure 7] individual devices are controlled by low level agents co-located with the devices they are controlling. Coordination between low-level devices is performed through direct supervision by a central software agent.

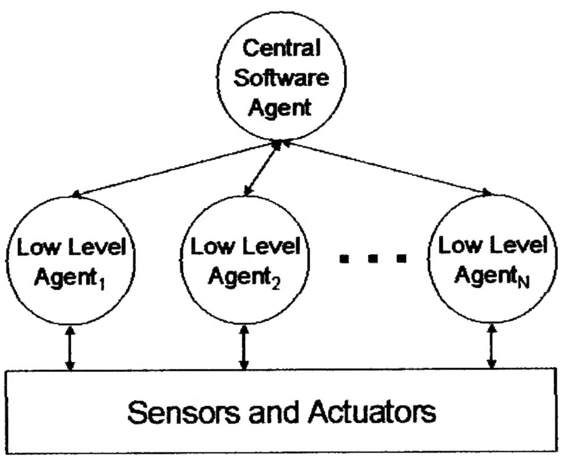

Figure 7 – Centralized Agent Topology

A variation of the centralized control theme is hierarchically organized control agents. [Figure 8] Hierarchical structures have successive tiers of agents that use direct supervision to control agents at a lower level and provide coordination between agents at the lower level. At the highest level, a single software agent provides coordination across the entire ship. Hierarchical agents allow designers to limit the scope of individual agents to manageable size by separation (dividing a system into multiple parallel components and by abstraction (reducing the fidelity with which an agent models the system it controls).

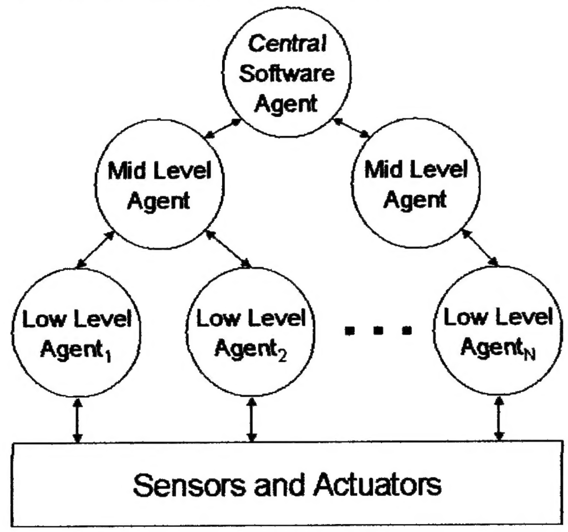

Figure 8 – Hierarchical Agent Topology

Software agents using mutual adjustment as a basis of coordination may be used to form flat, heterarchical structures. These agents use peer-to-peer relationships to satisfy ship-wide needs. [Figure 9] Coordination may be explicit, in which adjustments are mutually agreed upon by corroborating agents; or implicit, in which agents react to a common set of knowledge without explicit knowledge of other agents intentions. Regardless, heterarchical agents share knowledge of the underlying system as well as system goals.

The survivability of an autonomous system is heavily dependent upon the continued operation of its control system. In turn, the survivability of the control system is dependent upon the continued operation of the infrastructure upon which the control system operates. We can infer that the survivability of the hardware controllers, the agents that operate on the hardware controllers and the control network by which they communicate directly impact the survivability of the ships auxiliary systems. If the ship is to continue to function in the presence of battle damage, the control

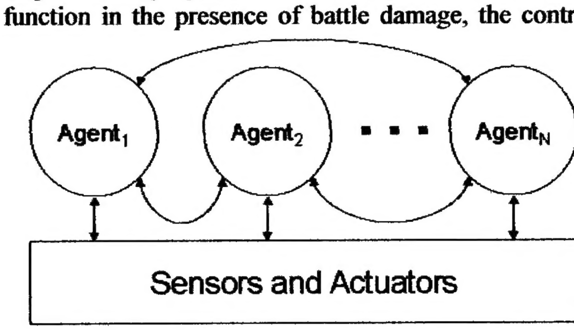

Figure 9 – Heterarchical Agent Topology

system must maintain the ability to reconfigure itself when processors, agents and network components fail. Centralized and hierarchical control systems, because they are dependent upon a single top level software agent for ship-wide coordination, are inherently less

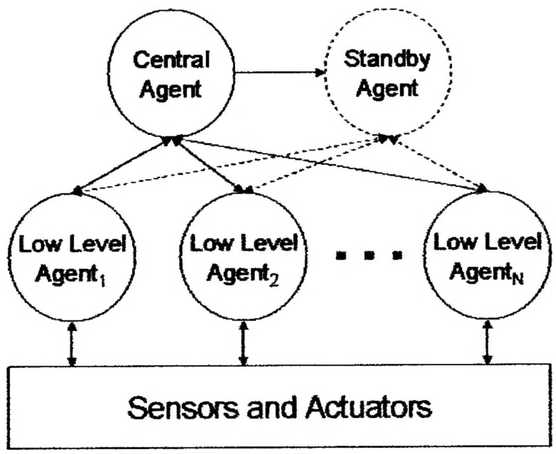

Figure 10 – Standby Agent

survivable than heterarchical systems. Robustness may be added to centralized and hierarchical systems by adding redundant processors hosting "standby" agents, albeit at the cost of additional infrastructure [Figure 10]. These standby agents operate passively, gathering data and maintaining an awareness of the system state but exerting no control over the physical system until the currently active agent fails. At this time, the standby agent becomes the active agent asserting control over the system.

VI. EMERGING APPLICATIONS

Intelligent Control, Distributed Intelligent Control and Software Agents are active bodies of research in the control systems and computer science communities. To date, little work has been done to transition these emerging techniques into ship control systems. A few representative examples of the work that has been performed are described here. Specifically, we describe the following experimental systems: the Smart Valve, an example of CLI; Starfish, a self-reconfiguring control network; the Open Autonomy Kernel, a framework for employing distributed intelligent agents; and the Integrated Engineering Plant.

Smart Valve

Smart valves are self-controlling valves that use sensed or inferred information on the valve (actuator) position, fluid flow rate, surrounding pressure and fluid temperature to diagnose the current state of the valve and proximate components [Figure 11]. By itself the smart valve is a useful tool for performing fault isolation, such as closing the valve when a pressure drop indicates a rupture within the system. Smart valves incorporate an embedded communications infrastructure that allows them to interface with a device or field-level network. Additionally, smart valves are programmable, allowing them to be adapted to future distributed intelligent control architectures. Rapidly maturing, smart valve technology was tested under live fire conditions [13], [14].

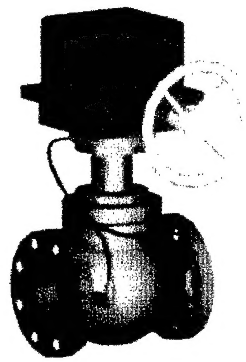

Figure 11 – Smart Valve

Starfish

Component Level Intelligence and Network Fragment Healing technology for reconfiguration was demonstrated on the ONR Afloat Lab, the YP-679 (the STARFISH) [Figure 12]. The YP demonstration addressed two of the three requirements for survivable automation: (1) Survivable Information Processing, and (2) Survivable Network Communications.

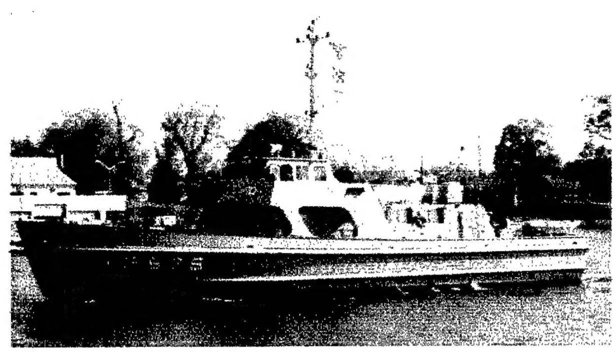

Fig 12 – YP-679 STARFISH

The YP Starfish project involved the development and demonstration of concepts for autonomous healing of component level shipboard networks utilized for automation of machinery systems. [15]. Automation system architecture was applied to the Propulsion Engines, the Fuel System, the Seawater Cooling

System, the Diesel Portion of the diesel generators, and the Steering System. The Afloat Lab automation system includes 183 sensors and actuators that are controlled and monitored by 85 computing nodes grouped into 8 different special locations throughout the vessel. By distributing the intelligence of the entire system towards individual components, each of the computing nodes is responsible for a relatively small portion of the overall control and monitoring function. If a processor fails, only a small part of the system's functionality is lost. Loss of a critical node or sensor could result in the inability to execute a needed control algorithm. Redundant sensors or actuators can be utilized, or data fusion using inputs from other independent sensors could be used, so that a new level of survivability can be reached.

The YP Starfish project showed that system reconfiguration using distributed intelligence can be achieved, with resultant continuity of platform services. Research issues remain, the most notable being scalability of technology to a full sized Naval platform. Other issues are associated with the reconfigurable electrical system capability to provide the same degree of power survivability that the YP communications system now possesses.

Open Autonomy Kernel

A prototype Intelligent Agent system called the Open Autonomy Kernel (OAK), designed to control auxiliary systems distributed on U.S. Navy surface combatants, has been developed and tested [16]. As an architecture for autonomous distributed control, OAK addresses control as a three-step process: diagnosis, planning and execution. OAK is specifically designed to support "hard" control problems in which the system is complex, sensor coverage is incomplete, and distribution of control is desired. A unique combination of model-based reasoning and autonomous agents are used. Model-based reasoning is used to perform diagnosis. Observations and execution are distributed using autonomous intelligent agents. Planning is performed with simple script or graph-spanning planners. OAK has been tested on a hardware simulation of the chilled water system of an ARLEIGH BURKE (DDG-51) Class destroyer. The simulator, known as the Chilled Water Reduced Scale Advanced Demonstrator (CW-RSAD) [Figure 13], consists of approximately 45 interdependent subsystems that must be coordinated by a higher-level control system in order to achieve a desired operational state. OAK has demonstrated control of system level responses (i.e. reconfiguration to meet mission goals) using the model based reasoning engine and an expert system implemented on a distributed device level control system to support normal operations.

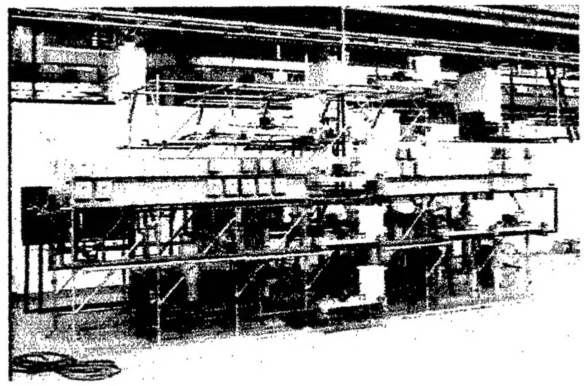

Fig. 13 – Chilled Water Test Bed

Integrated Engineering Plant

The Integrated Engineering Plant (IEP) is a dynamically reconfigurable and scaleable automation system architecture [Figure 14]. The overall goal is to progress from centralized control for each major ship system (propulsion, power generation, power distribution, damage control, weapons, etc) to an integrated automation system that provides for all aspects of ship control under varying conditions. The objective is to automate the engineering plant while inheriting all systems within a single reconfigurable network. This will be accomplished by developing and integrating three layers:

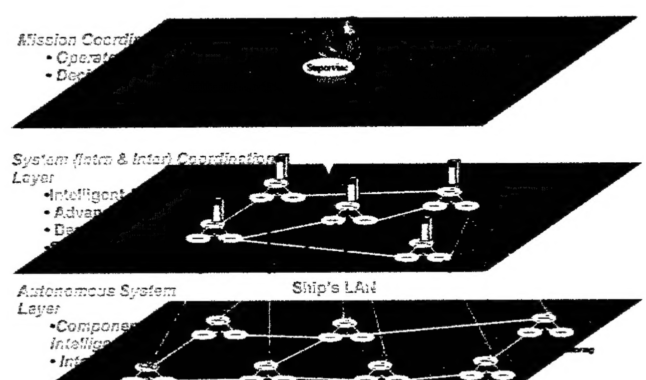

Figure 14 – Integrated Engineering Plant

1. A Module/Component Layer in which Component Level Intelligence will reside. Each module/component will be a node that communicates with the Shipboard Local Area

Network (LAN). Each node will have local computational capability. In addition, information and control for each module/component is resident in the node as well as in the LAN. Examples of nodes at this level are the propulsion module, compressors, converters and power distribution modules. Diagnostics are performed in this layer, and HMI functionality may be required.

2. A Process Layer concerned with diagnostics and with optimization of overall system performance by ensuring availability of any required system level function. The controllers at this level will communicate with system function controllers. Examples of system functions are propulsion, power generation, and damage control. HMI functionality is not required at this level.
3. A Mission Control Layer that is associated with ship wide resource allocation and planning. Management of all IEP systems are based on current ship status regarding mission, operating scenarios, priorities and available resources. [15]

VII. CONCLUSION

The complexity of ship control systems is increasing exponentially with each successive generation of ships. State-of-practice automated controllers will be incapable of effectively controlling next generation ships without exponentially increasing the workload of the individual sailor. Distributed intelligent control promises to advance automated control and enable control of complex ship systems with limited manning. Specifically, distributed intelligent control architectures will provide benefits with respect to the requirements for (1) enhanced recoverability and continuity of service to assure fight through capability through reconfiguration and resource reallocation on the fly, (2) reduced workload, resulting in lower manpower costs and Total Ownership Cost, and (3) reduced maintenance due to improved diagnostic capability and capability to predict impending failures.

[1] US NAVY SURVIVABILITY DESIGN HANDBOOK FOR SURFACE SHIPS Sept 2000 Chief of Naval Operations Ship Safety and Survivability Office N86DC The Pentagon

[2] B. Kuo, *Automatic Control Systems*, 4th ed., Prentice-Hall Inc. Englewood Cliffs, NJ 07632, (1982) pp 1-4.

[3] L. Fausett, *Fundamentals of Neural Networks*, Prentice Hall; 1st edition (1994).

[4] C. Watkins, Learning from Delayed Rewards, Thesis, University of Cambidge, England (1989).

[5] E. H. Shortliffe, *Computer-Based Medical Consultation: MYCIN*. Elsevier North, Holland (1976).

[6] F. Jensen, *An Introduction to Bayesian Networks*, UCL Press, London (1996).

[7] J. Kolodner, *Case-Based Reasoning*. San Francisco, California: Morgan Kaufmann (1993).

[8] A. Coddington, and R. Aylett, Plan generation for multiple autonomous agents: an evaluation. In Proc. 15th Workshop of the UK Planning and Scheduling SIG, Liverpool, Vol 1, 126-137, (1996).

[9] B. Kuipers: Reasoning with Qualitative Models. Artif. Intell. 59 (1-2): 125-132 (1993)

[10] T. Hentzinger, The Theory of Hybrid Systems in *Verification of Digital and Hybrid Systems* (M.K. Inan, R.P. Kurshan, eds.), NATO ASI Series F: Computer and Systems Sciences, Vol. 170, Springer-Verlag, 2000, pp. 265-292.

[11] Stuart Russell and Peter Norvig, *Solution Manual for "Artificial Intelligence: A Modern Approach."* Englewood Cliffs, NJ: Prentice Hall; 1995.

[12] H. Mintzbeg, *The Structuring of Organizations*. Prentice-Hall, Englewoods Cliffs, NJ (1979).

[13] S. Dibona and B. Callahan, "Development of a Flow Estimation Model for a Multi-mode Smart Valve," Proceedings of the 13th Ship Control Systems Symposium, April 7-9, Orlando, FL, USA (2003).

[14] "Real Time Detection of Damage to Shipboard Fluid Networks Using Enhanced Smart Valve Technology" presentation, Fairmount Automation, 25 Sept 2003

[15] Integrated Engineering Plant for Future Naval Combatants- Technology Assessment and Demonstration Roadmap APPENDIX: Contributing Team Documents MSD-50-TR-2003/01 January 2003

[16] Scheidt et al, Autonomous Distributed Intelligent Agents Using Model Based Reasoning, Proceedings of the 13th Ship Control Systems Symposium, Orlando, FL, USA (2003).

David H. Scheidt received a B.S. in Computer Engineering from Case Western Reserve University in 1985.
He is a group scientist at the Johns Hopkins university Applied Physics Laboratory. His current work concentrates on the research and development of distributed intelligent control systems. Previous efforts led by Mr. Scheidt include development of a meta-database containing ~1,000 heterogeneous databases, distributed immunology tracking, locomotive control, railroad dispatching, and distributed multi-level secure information systems. Mr. Scheidt has twice received the National Performance Review's Hammer Award. His email address is david.scheidt@jhuapl.edu.

Katherine F. Drew has been a Science and Technology Program Manager at the Office of Naval Research, Hull, Mechanical and Electrical Division of the Engineering, Materials and Physical Sciences Department since 2001. Prior to that time, while at ONR, she directed a university, industry and laboratory based research program in methods for affordable and robust product design. Before coming to ONR in 1994, she was Head of the Underwater Systems Analysis Branch at Naval Surface Warfare Center, White Oak. During the mid-1980's she worked in the private sector as a Principal Engineer in the area of digital signal processing. She received a B.S. in Physics from the University of Pittsburgh, and an MS in Applied Mathematics from the University of Maryland.